The Laws of Vulnerabilities 2.0

Black Hat 2009 Edition

Presented by
Wolfgang Kandek, CTO, Qualys, Inc.
http://laws.qualys.com
July 28, 2009

Abstract

This study of dynamics in the vulnerability life cycle began in 2001 when Qualys launched its global vulnerability scanning software-as-a-service called QualysGuard. In 2004 Qualys used accumulated scanning data to identify the "Laws of Vulnerabilities" – four distinct, quantifiable attributes used to drive strategies for protecting networks, systems and data. In this paper, Qualys re-examines these Laws based on a much larger anonymous sampling of data from 104 million vulnerability scans made in 2008 finding about 680 million vulnerabilities. For each of the Laws, the new data reveal:

Half-life:	Time interval for reducing occurrence of a vulnerability by half. Average duration of half-life continues to be about 30 days, varying by industry sector.
Prevalence:	Measures the turnover rate of vulnerabilities in the "Top 20" list during a year. Prevalence has increased, with 60% remaining in the list in 2008 compared to 50% in 2004.
Persistence:	Total life span of vulnerabilities. Persistence remains virtually unlimited.
Exploitation:	Time interval between an exploit announcement and the first attack. Exploitation is faster, often happening in less than 10 days compared to 60 days in 2004.

Wolfang Kandek is Chief Technology Officer at Qualys; email to wkandek@qualys.com.

Introduction

The near-universal availability of Internet connectivity has brought fundamental change in the way we use computer technology. Computers are now tools that are deeply integrated into our daily lives. They have become mobile and many people mix/perform both work and personal activities on their business computers, often operating outside of any corporate protection mechanisms. As such, these devices are our helpers – but they also are risky points of attack that can trigger theft of sensitive corporate or personal data. Understanding these risks is a prerequisite to taking steps that block potential attacks and keep networks, data and systems safe from harm.

Every moment they are connected to the Internet, networked computers are exposed to a hostile environment with thousands of threats probing every possible way to attack. This connectivity is a global conduit to attacks on weaknesses in operating systems, web pages serving viruses designed to exploit browser vulnerabilities, e-mails with malicious attached documents, and instant messaging chats with rigged pictures or video streams. Malware is sneaky, and has recently started to abuse the increasing awareness of security problems and to pose as security software. In those cases, attackers dupe users by charging them for the "privilege" of installing a useless and ultimately malicious application.

In the global open environment of Internet connectivity, it is crucial to assure the robustness of all involved computer systems, both servers and clients. And when vulnerabilities exist, IT administrators must ensure that systems are patched to create resilience and safety. This paper presents an analysis of actual global data that accurately define the dynamics of the vulnerability life cycle. It describes the key characteristics of half-life, prevalence, persistence and exploitation. These are the "Laws of Vulnerabilities" because they reliably describe what you can expect of a typical vulnerability. Understanding these attributes can help IT and security administrators prioritize efforts to fix critical vulnerabilities.

Contents

Half-life is the time interval measuring a reduction of a vulnerability's occurrence by half. Over time, this metric shows how successful efforts have been to eradicate a vulnerability. A shorter half-life indicates faster remediation.

Prevalence notes the turnover rate of vulnerabilities in the "Top 20" list during a year. Prevalent vulnerabilities are dangerous because they represent ongoing potent risks to computing environments. Risk rises as the prevalence rate rises because of the larger total number of top 20 risks tallied during a year.

Persistence measures the total life span of vulnerabilities. The fact that vulnerabilities persist and do not conclusively die off is a red flag for security administrators. It underscores the importance of patching all systems, and ensuring that old vulnerabilities are not inadvertently installed on new systems.

Exploitation is the time interval between an exploit announcement and the first attack. This metric indicates how much reaction time you might get before someone figures out how to exploit the vulnerability. The worst scenario is a "zero day" attack because there is no reaction time.

Data Sources

Data for this study comes from the QualysGuard KnowledgeBase, which includes a daily summary of all vulnerabilities found in scans executed by Qualys customers on their global network infrastructures. The scan data covers a six-year period; this paper focuses on the recent set from 2008. It includes 104 million vulnerability scans done that year with QualysGuard. Of those, 82 million were executed on internal Scanner Appliances, and 22 million from external Internet-based scanners.

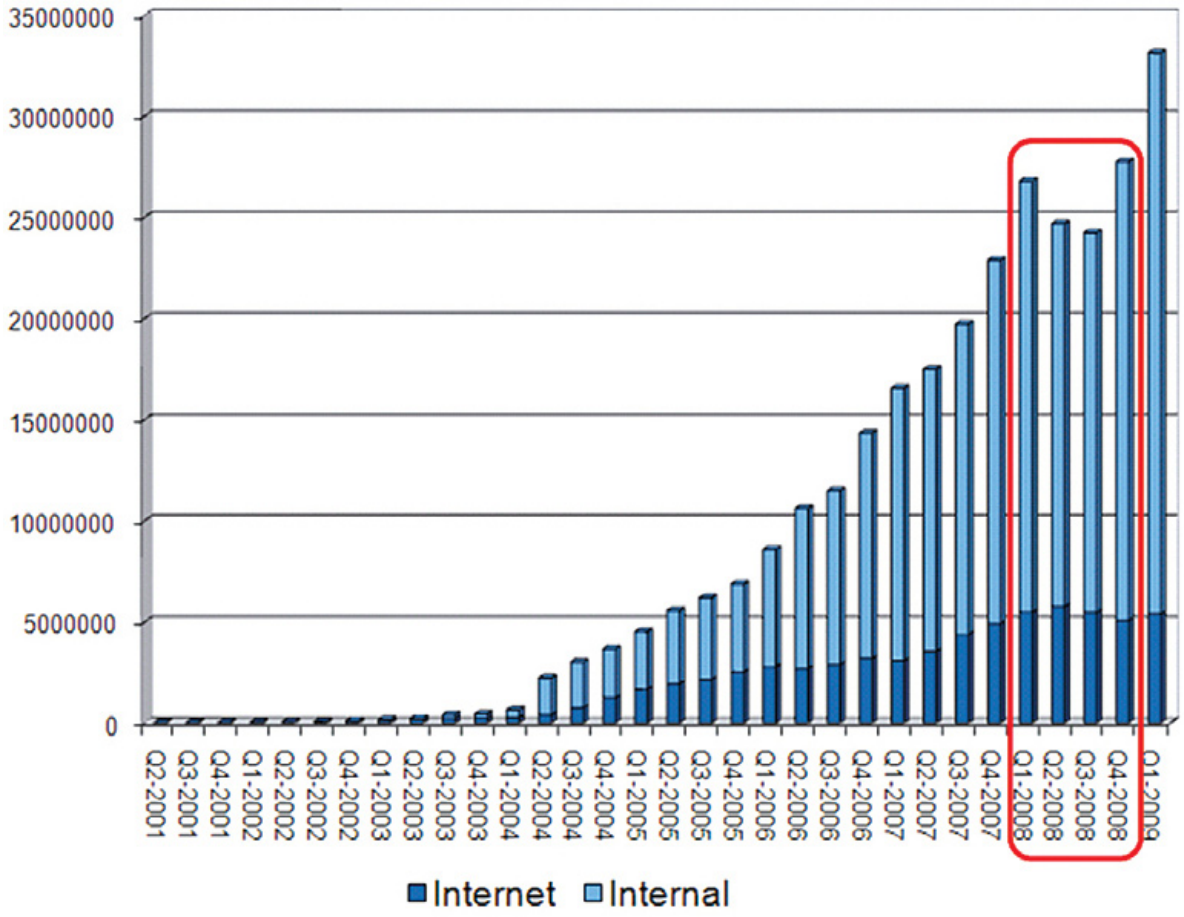

Figure D: Total Number of QualysGuard Scans (Internet & Internal)

The large number of scans made by about 3,500 organizations worldwide enables anonymous sampling, which protects the privacy of those customers. Data analysis includes all vulnerabilities for internal and external scans by the major industry sectors of Financial, Health, Manufacturing, Services and Wholesale/Retail.

Total vulnerabilities detected were close to 680 million, with more than 72 million having a severity of “critical.” This classification means a successful exploit will give the hacker control over the system.

Half–Life

The half-life of a vulnerability is the time required to cut its occurrence by 50 percent. This measure indicates how quickly IT administrators remediate vulnerabilities.

Of 72 million critical vulnerabilities, the data analysis shows that duration of half-life is now 29.5 days. The majority of vulnerabilities are now found in client-side applications, as reflected by the growth in scans done with internal Scanner Appliances.

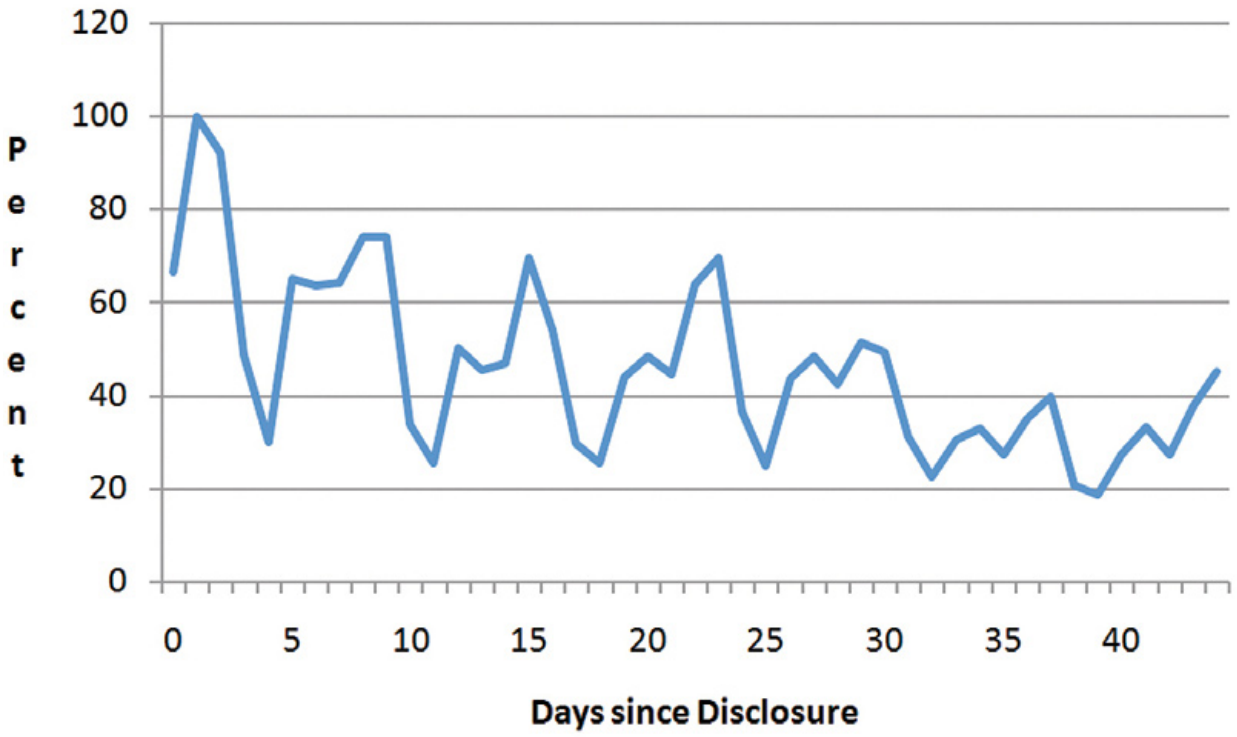

Figure H.1: Half-life for critical vulnerabilities during 2008

Figure H.1 shows that on average, IT administrators at Qualys customers take roughly 30 days to remediate critical vulnerabilities on half of their vulnerable workstations and servers. This does not represent any obvious improvement to scan data from 2004, which also was approximately 30 days. However, many of the factors affecting half-life have changed in the last four years so a direct comparison is difficult. In 2004, the number of disclosed vulnerabilities was less than half of the vulnerabilities found in 2008. Moreover, much of the research focus has changed from server-side vulnerabilities to vulnerabilities on the desktop. These span a much larger group of applications, including many third party programs, such as Adobe Reader, Apple QuickTime and other rich-media applications. At the same time, tools to manage the vulnerability and patch cycle have become more mature and have enabled IT administrators to implement automated mechanisms to apply patches and survey systems with a minimum of disruption.

The data also show differences in the way industry sectors execute remediation strategies. Service, Finance and Wholesale/Retail sectors are the most successful, posting vulnerability half-lives of 21, 23 and 24 days respectively.

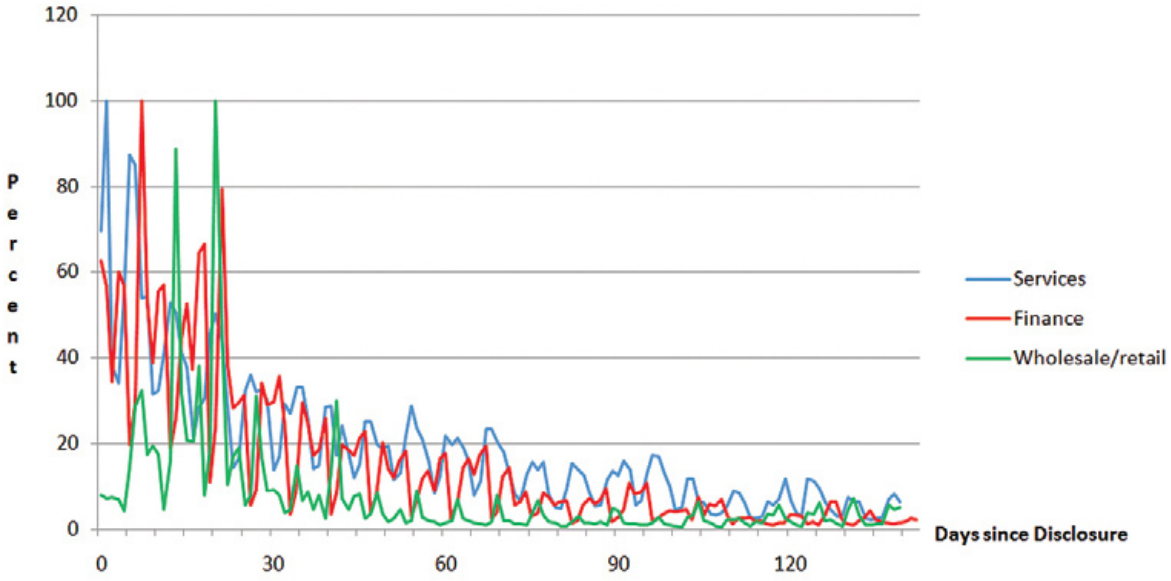

Figure H.2: Half-lives for Services, Finance and Wholesale/Retail

While Services, Finance and Wholesale/Retail are better than average, Health and Manufacturing are below the average with 38 days and 51 days respectively.

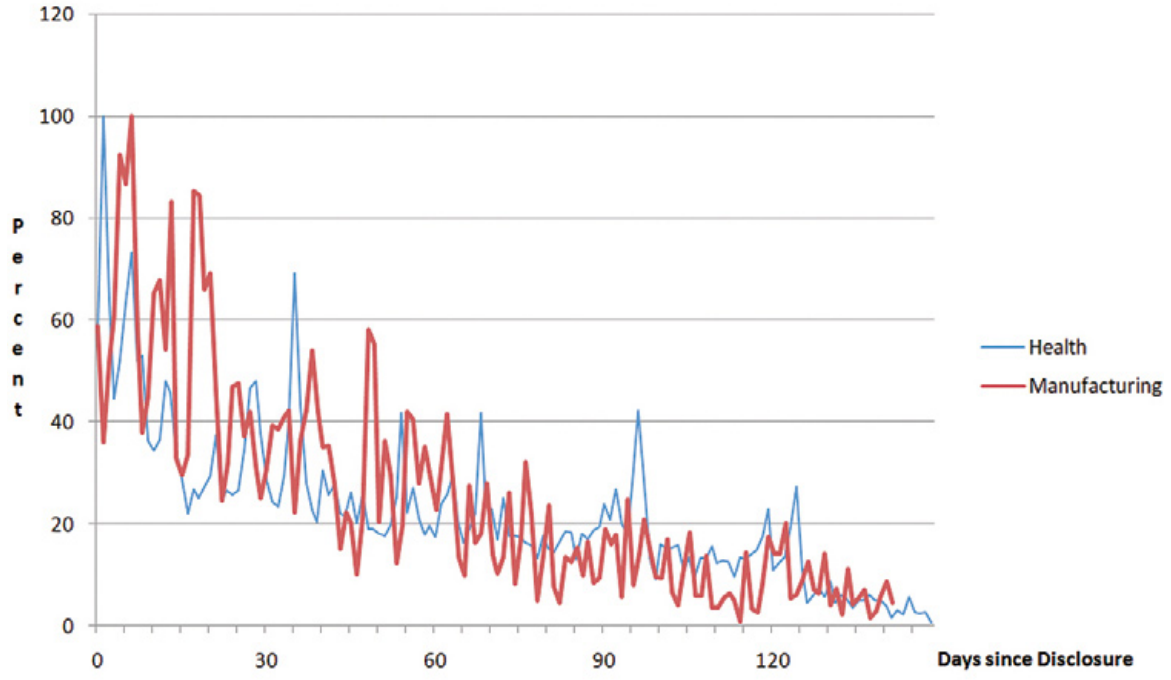

Figure H.3: Half-lives for Health and Manufacturing

Where the Data Came From

Data for this study came from vulnerability scans performed by Qualys customers in five major types of business sectors. Figures H.2 and H.3 refer to the following from the North American Industry Classification System (NAICS), which is the standard used to track and analyze statistics about the U.S. business economy. NAICS replaces the older Standard Industrial Classification (SIC) system. For more information, see **http://www.census.gov/eos/www/naics/index.html**.

Services	constitute the largest segment of the business economy. This sector includes Professional, Scientific, and Technical Services (NAICS Classification 54), Management of Companies and Enterprises (55), Administrative and Support and Waste Management and Remediation Services (56), Educational Services (61), Arts, Entertainment, and Recreation (71), Accommodation and Food Services (72), and Other Services except Publication Administration (81).
Finance and Insurance (52)	are firms that do financial transactions. These can involve creating, liquidating, or changing ownership of financial assets, or facilitating the transactions.
Wholesale Trade (42)	is generally involved with selling merchandise. This sector includes the production from agriculture, mining, manufacturing, and other industries such as publishing.
Health Care and Social Assistance (62)	are providers of health care and social assistance for individuals.
Manufacturing (31-33)	includes companies that make new products by mechanically, physically, or chemically transforming raw materials, substances or components.

Prevalence

Prevalence is the measure of non-renewal of detected vulnerabilities. To calculate prevalence, changes were noted in the Top 20 critical vulnerabilities for each month of 2008. At year end, prevalence was calculated as a percent value with each 5 percent unit indicating the substitution of one vulnerability in the Top 20. A turnover value of 20 – percent indicates that four vulnerabilities were substituted, 100 percent indicates that all vulnerabilities were substituted, and values above 100 percent mean that all vulnerabilities were substituted more than once in the Top 20.

The prevalence metric for 2008 was 60 percent, which means 12 vulnerabilities were substituted from the Top 20. It also means that eight critical vulnerabilities retained a constant presence from January of 2008 to December of 2008, which makes them attractive targets for attackers.

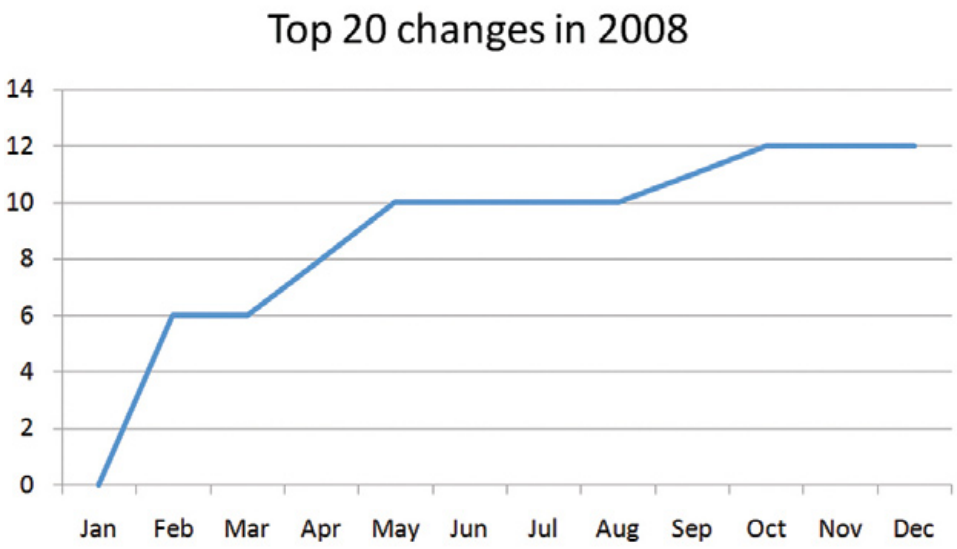

Figure P.1: Top 20 substitutions in 2008

So while it is interesting to look at the make-up of the new vulnerabilities, it is quite valuable to look at the applications that retained a constant presence in the Top 20:

- **Microsoft Office**
- **Windows 2003 Server SP2**
- **Sun Java**
- **Adobe Reader**

This list proves that there are applications that are not receiving enough attention by IT administrators. We can only speculate why certain applications are not being updated as aggressively as the core operating systems, but attackers have already picked up on these new vectors. Microsoft's Security Intelligence Report 2008 (Volume 6, p. 47ff) reports Adobe Flash, Apple QuickTime, RealPlayer, Microsoft Office and Adobe Reader vulnerabilities in their file format section, but does not disclose how they compare overall. F-Secure published two data for 2008 and 2009 showing an increase in attacks targeting PDF vulnerabilities:

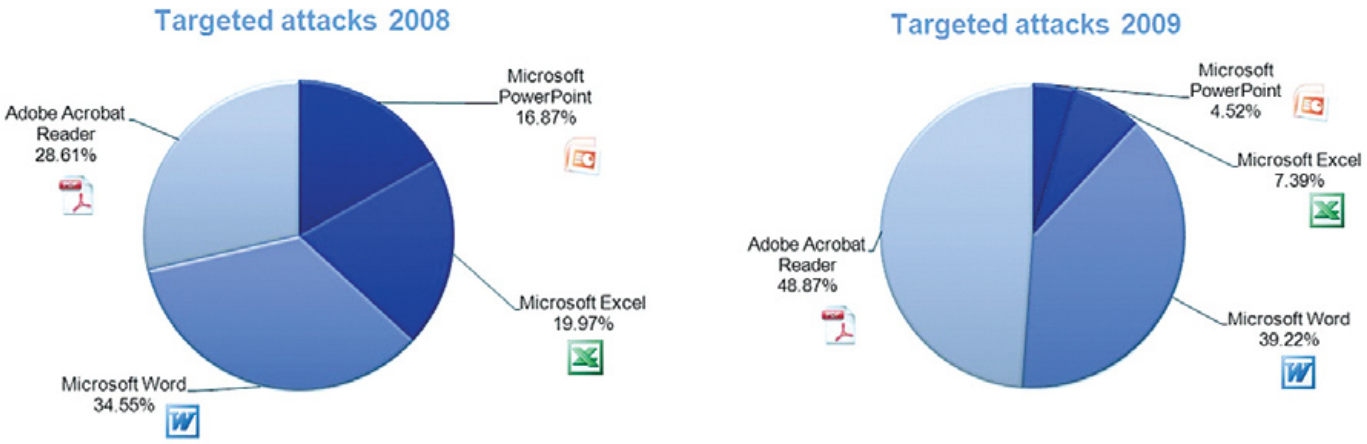

Figure P.2: Changes in attack formats for Documents

More clarity on remediation speed for some of these vulnerabilities comes from a 2009 dataset. Patches for critical Adobe Reader vulnerabilities APSA09-01 and APSA09-02 were published in the middle of March 2009 and the middle of May 2009, respectively. Plotting the detection rate of these vulnerabilities indicates how much of the installed base of Adobe Reader was fixed before release of the second patch set. In an ideal scenario, most Adobe Readers would be patched for March's APSA09-01 by the middle of May – about two months after the patches became available. However, the data show that at maximum, the delta between the two vulnerability occurrences is 20 percent. More than 80 percent of all systems were still in need of both patches, clearly a sign that too few systems had been patched in the preceding two months with March's APSA09-01.

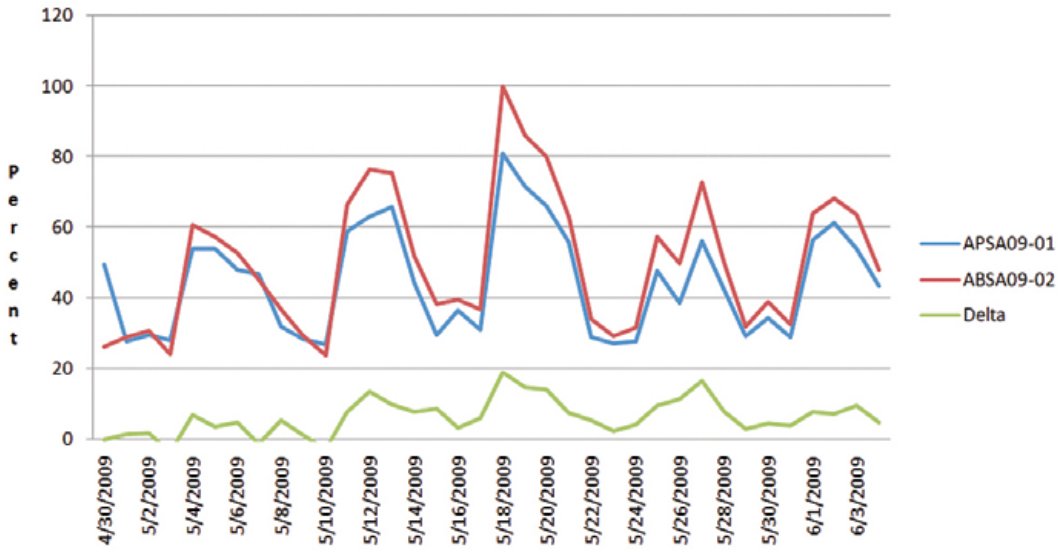

Figure P.3: Comparing Adobe Patch level over time

Persistence

Persistence is the measure of longevity for a vulnerability. Initially, research expectations were for persistence to be under the one percent margin after seven months, based on the half-life measure described above. Scan data do not confirm this expectation. All vulnerabilities, even the critical ones, stabilize at the 5-10% level. As an example, the following graph tracks critical Microsoft vulnerabilities published in January, February, March and April of 2008 in areas such as Core Operating system, WebDAV, Outlook and Graphics libraries – all of them showing similar behavior with stabilization around the seven percent level.

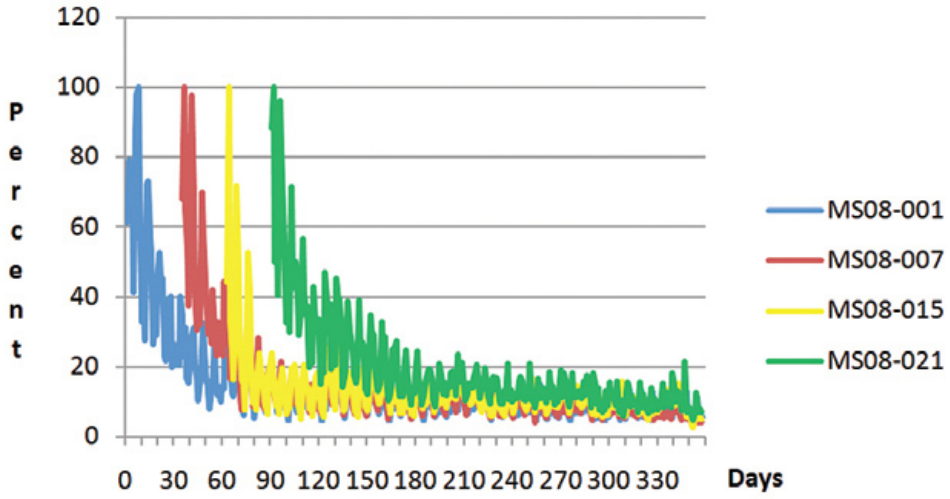

Figure P.1: Persistence levels of critical Microsoft vulnerabilities of 2008

Older vulnerabilities continue to persist today, such as MS02-039/061, an announcement that contained the patches for a weakness in MS-SQL Server used by the SLAMMER worm in 2003. TippingPoint's ThreatLinQ sensors network still reports more than 70,000 Slammer-infected machines – and indicates even new infections:

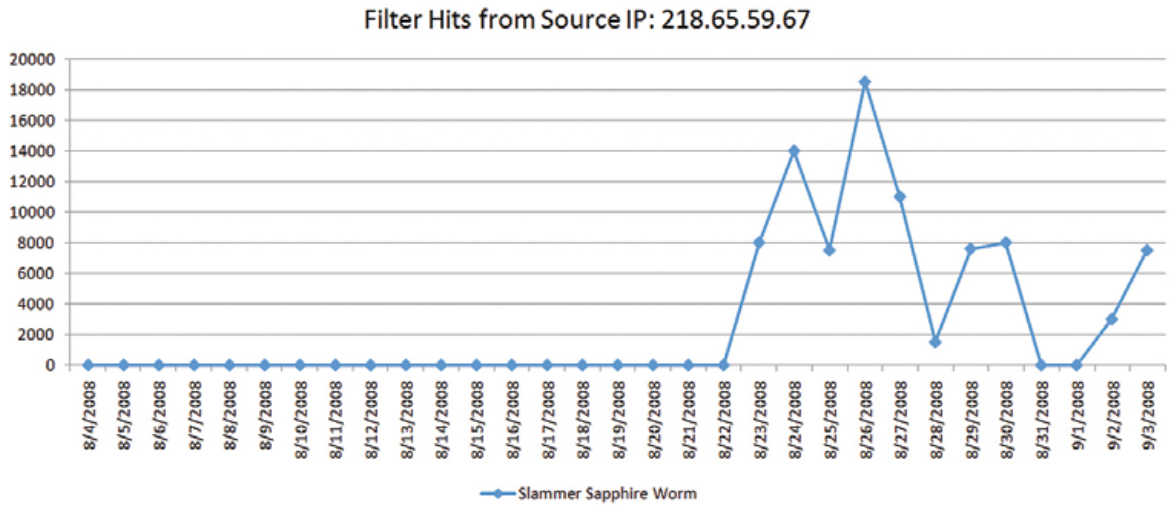

Figure P.3: Example system infected in late 2008 with Slammer

Data from this study confirms the existence of a baseline level that has been constant for the last year. This persistence is astounding for such a highly visible and ancient vulnerability such as MS02-Slammer.

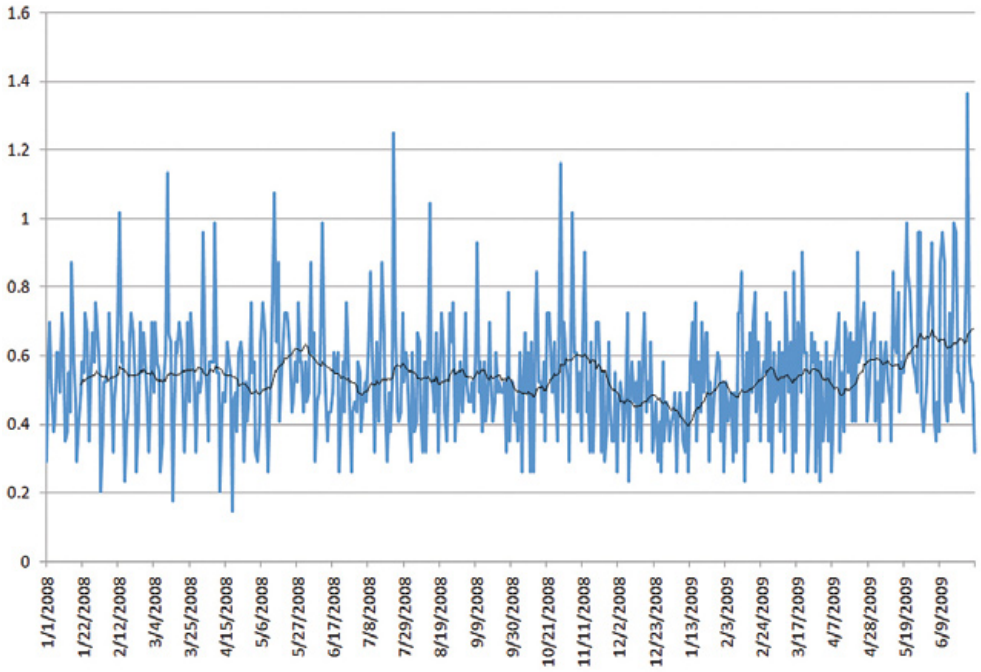

Figure P.3: Comparing Adobe Patch level over time

In the best case, these are systems that are probably beyond the reach of patching due to legacy application requirements, so protection usually relies on mitigating technologies such as Intrusion Prevention Systems. Often, infection by old vulnerabilities occurs when new systems are deployed with old builds of software that contain and reintroduce these vulnerabilities into the networks. In the worst case, these systems are simply forgotten, and security administrators have yet to discover their vulnerabilities.

Exploitation

Exploitation is the measure for the time needed for an attacker to craft and publish an exploit for a known vulnerability. Research data in 2004 showed 80 percent of critical vulnerabilities had an exploit available within 60 days following their release. In 2008 and 2009, this time has decreased significantly and many exploits are becoming available in less than 10 days. Today, the expression “zero-day” is well known in IT administration and security circles, defined as having exploits in the wild when the vulnerability is made public. This type of exploit leaves IT administrators with no time to execute their procedures for vulnerability intelligence, patch evaluation and deployment.

QualysGuard currently notes 56 zero-day vulnerabilities of varying criticality in its KnowledgeBase, and companies such as VeriSign/iDefense and TippingPoint typically have hundreds of advisories open with vendors that await corresponding patches and publication. In 2008 Microsoft was forced twice to release patches outside of their normal release schedule; in 2009 so far three times, each time acknowledging exploit usage was significant enough to go through this ad-hoc process. Microsoft’s April 2009 security bulletin had 21 vulnerabilities, 10 of which noted indications that an exploit was either available or being worked on. Zero-day vulnerabilities are by no means limited to Microsoft; they span the entire industry. Adobe had two zero-day instances in 2009 as of this study’s publication; both exploits attacked the PDF file format. Apple had several zero-day vulnerabilities in Mac OS X, most recently in the Java component. In 2008 both iTunes and QuickTime were affected. Popular open source browser Firefox had three zero-day vulnerabilities in 2008.

The increased availability of exploits ultimately means that IT administrators have to protect their systems in a shorter time. The easiest way to accomplish this is to apply patches as soon as they are released and thereby to lower the window of exposure. Remediation may not be possible on all systems, such as those with criticality or uptime constraints. Consequently, IT administrators should segment their systems into fast patch and slow patch pools and review the possibility of using mitigating technologies for protecting the slower patch pools. The existence of a fast patch pool has the additional benefit of gathering additional in-house experience related to the application of the patch and its potential side effects.

Another possibility of shrinking the exposure window is to segment applications into fast patch and slow patch segments, putting high exposure programs such as browsers and their plug-ins into the fast segment. Browser companies Google with Chrome and Mozilla with Firefox have integrated their own update mechanisms that use quite aggressive update rhythms. Recent research has shown the effectiveness of these approaches, where over 90 percent of all Chrome browsers and over 80 percent of all Firefox browsers were updated to latest level within 10 days following the release. Other popular browsers such as Safari and Opera only reach the 40 percent or 20 percent mark, in the same timeframe.

Summary

The scan data analyzed in this study show that the challenges faced by IT administrators in securing their systems have grown in the last four years. There is near-universal Internet connectivity and increasing reliance on Internet-based applications to conduct modern business. Both trends multiply the attack surface that malware authors are using to gain control of target systems. Modern attacks are multi-pronged and probe weaknesses on multiple levels in a single download – starting with the operating system, going through browsers, their plug-ins and common installed applications such as Office suites and Media players.

There are steps that IT and security administrators can take to gain the upper hand in this battle. Visibility over the network is the important first step, which entails mapping the overall network and systems to get an accurate graph of connectivity between machine clusters. Next is to evaluate the patch level of the operating system of individual machines and to enumerate installed software – both are necessary to form a picture of the types of machines and applications. Information from the first two steps will help quickly isolate systems at risk. Then develop a patching strategy based on criticality of the machines in the network and segmentation of the applications, most likely using a tool for the automatic application of these patches.

Ubiquitous Internet connectivity of computer systems clearly amplifies the risk of security problems. Organizations could react with an “ostrich strategy” and filter all Internet access. In some cases, it could be useful to completely isolate certain machines from network access. But in most business environments, restricted Internet access simply won’t work. Many business users legitimately try new communication tools intended for consumers such as instant messaging, auctions, forums, social networking and Twitter – and integrate them with their workflow to enhance productivity. Locking down systems defeats the potential business utility of messaging and social media applications. Lockdowns also alienate users, which prompts them to develop workarounds or simply stop using those systems. Even the relatively mild security mechanisms imposed by Windows Vista have triggered negative feedback from the user community and slowed that operating system’s adoption cycle.

Modern computer systems need to provide their users with the tools for full productivity, including connectivity to the Internet and open access to external data sources. Correspondingly, systems must have a level of robustness that is high enough to block the constant attacks coming through the Internet. As validated by the Laws of Vulnerabilities, the fundamental requirement for robustness is to keep computer systems updated with current technology, including rapid deployment of new program versions and patches. Comparing the state of vulnerabilities in your organization’s network and systems to averages reflected in the Laws of Vulnerabilities will help gauge progress toward a more robust and secure environment.

Case Study: Conficker/Downadup - MS08-067

Security bulletin MS08-067 was published by Microsoft in October 2008. The vulnerability was remotely attackable and exploit code was available for purchase in the wild. Microsoft decided to announce it out of sync with its scheduled release on November's "Patch Tuesday" due to its severity. The security community largely agreed with that decision and assigned the vulnerability the highest criticality rating. About a month later, in November, the security industry recorded first signs of a new worm that coupled the vulnerability described in MS08-067 with the type of scanning and identification mechanism needed to reach fast-growth spreading factors. The worm was called Conficker or Downadup; its advent triggered a phase of intense activity for IT administrators. Security researchers were able to reverse engineer the workings and communication mechanisms of the worm and attested to its innovative features. The security industry formed the "Conficker Working Group" to collaborate on minimizing the worm's impact. The worm's creator responded with Conficker versions B, C, D and E, each engineered to circumvent these protection mechanisms and to improve the infection capabilities.

The worm reached many machines (estimates range from 3-15 million systems) and formed a large global botnet. More than 60 percent of infected systems are in the extended BRIC region (Brazil, Russia, India, China, etc.). Outbreaks have occurred in government, military, health and commercial networks.

By the end of March, security researchers in Germany found a way to remotely detect the presence of an infected machine. They alerted the vulnerability scanning industry, which released the scanning capability in a concerted, synchronized effort. Scanning for Conficker reached new heights – Qualys' scan volume sextupled as many customers swept their entire Windows infrastructure. The infection rate for these systems was under one percent, much lower than the six percent reported by other industry sources – an indication of successful security/patching mechanisms.

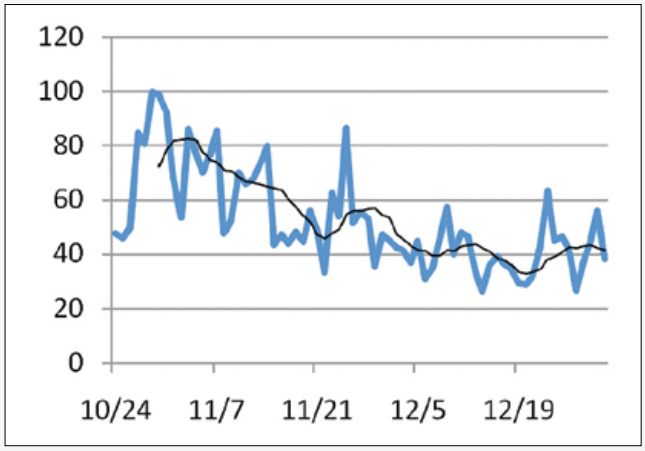

Figure C.1: MS08-067 half-life

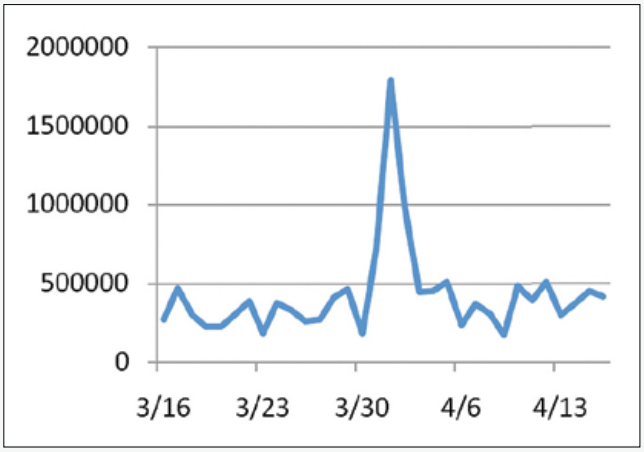

Figure C.2: IPs scanned

Nevertheless, today Conficker is alive and well, and its botmasters are renting out capabilities on exploited systems to serve malware and send spam. The total exact cost of the damage caused by the worm is unknown. In one case, the city of Manchester in the United Kingdom published an audit report that puts their cost of fixing the problems at roughly US$ 1.5 million.

References and Further Reading

The Laws of Vulnerabilities: Six Axioms for Understanding Risk, 2005
Gerhard Eschelbeck, Qualys
http://www.qualys.com/docs/Laws-Report.pdf

Why silent updates boost security, 2009
Stefan Frei, ETH Zurich
Thomas Duebendorfer, Google

Firefox (in)Security Update Dynamics Exposed, 2009
Stefan Frei, Bernhard Plattner, ETH Zurich
Thomas Duebendorfer, Google

F-Secure PDF file format vulnerabilities
http://www.f-secure.com/weblog/archives/00001676.html

TippingPoint ThreatLinQ – Slammer
http://dvlabs.tippingpoint.com/blog/2008/09/05/threatlinq-taking-out-the-trash

Conficker Working Group
http://www.confickerworkinggroup.org/wiki/

Manchester City Council Report on ICT service interruption, 2009
http://www.manchester.gov.uk/egov_downloads/Item_11.pdf

Like teenagers, computers are built to hook up, Cory Doctorow, 2009
http://www.guardian.co.uk/technology/2009/jun/16/computer-security-abstinence